Linear Optimization Duality and Sensitivity Analysis

Shixuan Zhang

ISEN 320-501, Fall 2023

1 Dual Linear Optimization Formulation

We now discuss *duality*, which is perhaps one of the most important concepts in mathematical optimization. The following example gives a motivation for our discussion.

Example 1. *A carpenter makes tables and chairs for sale. A total of 150 board ft of oak and 250 board ft of pine are available. A table requires 16 board ft of oak and 30 board ft of pine, while a chair requires 5 board ft of oak and 12 board ft of pine. Each table can be sold for \$40, and each chair for \$15. Assuming that the tables and chairs can be produced in fractions, we can define x_1,x_2 to be the number of tables and chairs to make, respectively, and write a LO model to maximize the revenue for the carpenter*

$$\begin{aligned} \max \quad & 40x_1 + 15x_2 \\ \text{s. t.} \quad & 16x_1 + 5x_2 \le 150, \quad \textit{(oak constraint)} \\ & 30x_1 + 12x_2 \le 250, \quad \textit{(pine constraint)} \\ & x_1, x_2 \ge 0. \end{aligned}$$

Now suppose the carpenter can sell any unused oak and pine to a wood company, and also buy additional oak and pine from it at the price of y_1 and y_2, respectively. How should the wood company set the prices such that the carpenter would have the same revenue as before? To answer this question, we can formulate this as a min-max problem

$$\min_{y_1,y_2\ge 0} \quad \max_{x_1,x_2\ge 0} \quad 40x_1 + 15x_2 + (150 - 16x_1 - 5x_2)y_1 + (250 - 30x_1 - 12x_2)y_2.$$

The min-max (or sometimes minimax*) problem lets the wood company determine the price first, and then let carpenter decides the production decisions, which imply the buying/selling strategies of the oak and pine. In game theory terminology, we are trying to find the* equilibrium *prices between the carpenter and the wood company. Note that the inner maximization problem* *can be written as*

$$\max_{x_1,x_2\geq 0} \quad 150y_1 + 250y_2 + (40 - 15y_1 - 30y_2)x_1 + (15 - 5y_1 - 12y_2)x_2.$$

If the either of the coefficients $40 - 15y_1 - 30y_2$ and $15 - 5y_1 - 12y_2$ is positive, then we can simply increase the value of x_1 or x_2 to make the objective value arbitrarily large. Thus we should have both coefficients being nonpositive. Consequently, we have an obvious solution to the inner maximization problem $x_1^* = x_2^* = 0$, so the min-max problem reduces to

$$\begin{aligned}
\min \quad & 150y_1 + 250y_2 \\
\text{s. t.} \quad & 15y_1 + 30y_2 \geq 40, \quad \textit{(nonpositive coefficient of } x_1\textit{)} \\
& 5y_1 + 12y_2 \geq 15, \quad \textit{(nonpositive coefficient of } x_2\textit{)} \\
& y_1, y_2 \geq 0.
\end{aligned}$$

This minimization problem is called the dual *problem to the original maximization problem. We will see below that solving the dual problem gives us prices that let the revenue of the carpenter stays the same as before.*

In general, we can formulate dual linear optimization (LO) problems based on the idea of *relaxation*. Consider a LO problem with some given $m' \leq m$ and $n' \leq n$:

$$\begin{aligned}
\min \quad & \sum_{i=1}^{n} c_i x_i \\
\text{s. t.} \quad & \sum_{i=1}^{n} a_{ji} x_i \geq b_j, \quad \forall\, j = 1, \ldots, m', \\
& \sum_{i=1}^{n} a_{ji} x_i = b_j, \quad \forall\, j = m'+1, \ldots, m, \\
& x_i \geq 0, \quad i = 1, \ldots, n', \\
& x_i \in \mathbb{R}, \quad i = n'+1, \ldots, n.
\end{aligned} \tag{1}$$

We can relax all the constraints (both equalities and inequalities) and penalize the violation with a price vector $y \in \mathbb{R}^m$ as

$$\min \quad \sum_{i=1}^{n} c_i x_i + \sum_{j=1}^{m} \left(b_j - \sum_{i=1}^{n} a_{ji} x_i \right) y_j. \tag{2}$$

The first part here is the original objective function, while the second part is given by the penalty on any constraint violation. Since we are minimizing the objective value, we would like to have a positive penalty $(b_j - \sum_{i=1}^{n} a_{ji} x_i) y_j > 0$ only when there is violation of this constraint, i.e., $\sum_{i=1}^{n} a_{ji} x_i < b_j$ for any $j \leq m'$, or $\sum_{i=1}^{n} a_{ji} x_i \neq b_j$ for any $j \geq m'+1$. Thus we should set y_j to be nonnegative for each $j \leq m'$ and unrestricted in sign for $j \geq m' + 1$. Now we want to find a *tightest* relaxation, in the sense that the objective value should be as close to the original objective value as possible. Naturally this leads to maximization in the variables $y_1, \ldots, y_m$. By rearranging terms in (2), we can write the tightest relaxation problem as

$$\max_{\substack{y_1,\ldots,y_{m'} \geq 0 \\ y_{m'+1},\ldots,y_m \in \mathbb{R}}} \quad \min_{\substack{x_1,\ldots,x_{n'} \geq 0 \\ x_{n'+1},\ldots,x_n \in \mathbb{R}}} \quad \sum_{i=1}^{n} \left(c_i - \sum_{j=1}^{m} a_{ji} y_j \right) x_i + \sum_{j=1}^{m} b_j y_j. \tag{3}$$

Note that if $c_i - \sum_{j=1}^{m} a_{ji} y_j < 0$ for some $i \leq n'$, then the inner minimization would be unbounded (as we can take x_i to be arbitrarily large). The same behavior happens if $c_i - \sum_{j=1}^{m} a_{ji} y_j \neq 0$ for some $i \geq n' + 1$, because x_i is not restricted in sign. Therefore, we can impose the constraints

$$\begin{aligned} \sum_{j=1}^{m} a_{ji} y_j &\leq c_i, \quad i = 1, \ldots, n', \\ \sum_{j=1}^{m} a_{ji} y_j &= c_i, \quad i = n'+1, \ldots, n, \end{aligned} \tag{4}$$

in the outer maximization without any compromise of its optimality. Consequently, the first summation in (3) vanishes and the inner minimization becomes trivial. We have therefore derived our *dual* problem, i.e., finding the tightest relaxation, as

$$\begin{aligned} \max \quad & \sum_{j=1}^{m} b_j y_j \\ \text{s.t.} \quad & \sum_{j=1}^{m} a_{ji} y_j \leq c_i, \quad i = 1, \ldots, n', \\ & \sum_{j=1}^{m} a_{ji} y_j = c_i, \quad i = n'+1, \ldots, n, \\ & y_j \geq 0, \quad j = 1, \ldots, m', \\ & y_j \in \mathbb{R}, \quad j = m'+1, \ldots, m. \end{aligned} \tag{5}$$

To distinguish the problems, we also call the original problem (1) the *primal* problem. The derivation above is known as *Lagrangian duality* and has been used beyond LO problems, e.g., in some nonlinear optimization or integer optimization problems as well. To save the effort of deriving the dual formulation from scratch each time, we summarize the correspondence of the sign restrictions and inequality directions for LO problems in Table 1.

Table 1: Dual correspondence of constraints and variables

Minimization		Maximization
$=$ constraint	$\longleftrightarrow$	free variable
$\geq$ constraint	$\longleftrightarrow$	nonnegative variable
$\leq$ constraint	$\longleftrightarrow$	nonpositive variable
free variable	$\longleftrightarrow$	$=$ constraint
nonnegative variable	$\longleftrightarrow$	$\leq$ constraint
nonpositive variable	$\longleftrightarrow$	$\geq$ constraint

2 Weak and Strong Duality

For notational convenience, we consider matrix forms of the primal and the dual problems:

$$\begin{aligned} \min \quad & c^\top x \\ \text{s.t.} \quad & Ax = b, \\ & x \geq 0, \end{aligned} \tag{P}$$

and

$$\begin{aligned} \max \quad & b^\top y \\ \text{s.t.} \quad & A^\top y \leq c, \\ & y \in \mathbb{R}^m. \end{aligned} \tag{D}$$

We can see that for any primal feasible solution $x = (x_1, \ldots, x_n)$ and dual feasible solution $y = (y_1, \ldots, y_m)$, we have

$$b^\top y = (Ax)^\top y = x^\top (A^\top y) \leq c^\top x. \tag{6}$$

Thus if the primal problem (P) admits an optimal objective value z^* with an optimal solution x^*, then

$$b^\top y \leq c^\top x^* = z^*, \tag{7}$$

for any dual feasible solution $y \in \mathbb{R}^m$, which means the dual problem (D) is bounded. Similarly, if the dual problem (D) admits an optimal objective value w^* with an optimal solution y^*, then

$$w^* = b^\top y^* \leq c^\top x, \tag{8}$$

for any primal feasible solution $x \geq 0$, which means the primal problem (P) is bounded. In the case where both the primal and the dual problems have optimal solutions, x^* and y^*, respectively, we have the inequality

$$w^* = b^\top y^* \leq c^\top x^* = z^*. \tag{9}$$

This is called the *weak duality* of LO problems. We may further extend our discussion to some infeasible or unbounded LO problems. Recall that we can set $z^* = +\infty$ (resp. $z^* = -\infty$) if the primal minimization problem (D) is infeasible (resp. unbounded), and $w^* = -\infty$ (resp. $w^* = +\infty$) if the dual maximization problem (D) is infeasible (resp. unbounded). If the primal problem is unbounded, then for any dual feasible solution y, there exists a primal feasible solution x such that

$$b^\top y > c^\top x,$$

which contradicts with the inequality (9). Thus the dual problem must be infeasible, in which case we may write $z^* = w^* = -\infty$. Alternatively, if the dual problem is unbounded, then by the same argument we must have an infeasible primal problem, in which case we may write $z^* = w^* = +\infty$. In fact, the equality $z^* = w^*$ holds generally for feasible primal and dual problems as well. This is known as the *strong duality* result for LO problems.

Proposition 1. *If the primal problem* (P) *has an optimal solution x^*, then the dual problem* (D) *has an optimal solution y^* with $c^\top x^* = b^\top y^*$.*

Proof. By assumption, the primal problem (P) is feasible and bounded. By the simplex method (with Bland's pivot rule), there exists a basic feasible solution $x = (x_B, x_N)$, for some basis $B \subset \{1, \dots, n\}$ and $N := \{1, \dots, n\} \setminus B$, such that $x_B = A_B^{-1} b$ and $x_N = 0$. Obviously, $c^\top x^* = c^\top x$ as both are optimal primal solutions. Let $y^* := (c_B^\top A_B^{-1})^\top$. We first check that y^* is a feasible solution to the dual problem (D):

$$A^\top y^* = \begin{bmatrix} A_B^\top \\ A_N^\top \end{bmatrix} (c_B^\top A_B^{-1})^\top = \begin{bmatrix} c_B \\ (c_B^\top A_B^{-1} A_N)^\top \end{bmatrix} \leq \begin{bmatrix} c_B \\ c_N \end{bmatrix} = c.$$

Here the last inequality is ensured by optimality of x, where the reduced cost vector $r = c_N - (c_B^\top A_B^{-1} A_N)^\top \geq 0$ (note that it is minimization). Then $b^\top y^* = c_B^\top A_B^{-1} b = c_B^\top x_B = c^\top x$, which implies that y^* is an optimal solution to the dual problem (D) by the weak duality (9). □

Using the symmetry between the primal and the dual problems, Proposition 1 tells us both have the same optimal value as long as one of them has an optimal solution. It remains to ask the possible outcomes of the primal and the dual problems if we know one of them is infeasible. The next example shows that it is possible for both of the problems to be infeasible.

Example 2. *Let*

$$A = \begin{bmatrix} 2 & -1 \\ -2 & 1 \end{bmatrix}, b = \begin{bmatrix} 2 \\ -3 \end{bmatrix}, c = \begin{bmatrix} -5 \\ 2 \end{bmatrix}.$$

Then the primal problem (P) becomes

$$\begin{aligned} \min \quad & -5x_1 + 2x_2 \\ \text{s. t.} \quad & 2x_1 - x_2 = 2, \\ & -2x_1 + x_2 = -3, \\ & x_1, x_2 \geq 0, \end{aligned}$$

which is obviously infeasible due to the conflicting constraints. The dual problem (D) can be written as

$$\begin{aligned} \max \quad & 2y_1 - 3y_2 \\ \text{s. t.} \quad & 2y_1 - 2y_2 \leq -5, \\ & -y_1 + y_2 \leq 2, \\ & y_1, y_2 \in \mathbb{R}, \end{aligned}$$

which is also infeasible because the second constraint implies $2y_1 - 2y_2 \geq -4$*, which contradicts with the first constraint.*

We summarize the possible outcomes of the primal and the dual problems in the following table, where the columns and the rows correspond to the primal and the dual LO problem, respectively, and a ✓ means possible while ✗ means impossible.

Table 2: Possible outcomes of the primal and the dual LO problems

	Optimal	Infeasible	Unbounded
Optimal	✓	✗	✗
Infeasible	✗	✓	✓
Unbounded	✗	✓	✗

A very useful implication of the strong duality (Proposition 1) is a relation between any primal and dual solutions, which is called *complementary slackness*.

Corollary 2. *Let x denote a feasible solution to the primal problem (P) and y a feasible solution to the dual problem (D). Then x and y are simultaneously optimal solutions to their LO problems if and only if the following condition hold:*

$$(c - A^\top y)^\top x = 0.$$

Proof. First, suppose x and y are optimal solutions. Then by Proposition 1, we know that

$$0 = c^\top x - b^\top y = (c - A^\top y)^\top x.$$

Conversely, by the same inequality we know that $c^\top x = b^\top y$. Now apply the weak duality (9), we know that both x and y are optimal solutions. □

Corollary 2 tells us the followings: for an optimal primal solution x^* and an optimal dual solution y^*,

(i) if $x_i^* > 0$ for some $i = 1, \ldots, n$, then we must have $\sum_{j=1}^m a_{ji} y_j^* = c_i$; or

(ii) if $\sum_{j=1}^m a_{ji} y_j^* < c_i$ for some $i = 1, \ldots, n$, then $x_i^* = 0$.

While our proof here is based on the particular format used in the primal (P) and the dual (D) problems, it is straightforward to check that the complementary slackness holds for any general LO forms. We may interpret it as the following simple rule:

> Either a variable is at its bound zero, or the corresponding dual constraint must hold as an equality.

We illustrate how we can use the complementary slackness to find an optimal solution of a pair of LO problems below.

Example 3. *Consider the LO problem:*

$$\begin{aligned} \min \quad & 3x_1 + 4x_2 + 2x_3 \\ \text{s.t.} \quad & x_1 + 2x_2 + x_3 = 5, \\ & 2x_1 + 3x_2 + x_3 = 8, \\ & x_1, x_2, x_3 \geq 0. \end{aligned}$$

The dual LO problem is

$$\begin{aligned} \max \quad & 5y_1 + 8y_2 \\ \text{s.t.} \quad & y_1 + 2y_2 \leq 3, \\ & 2y_1 + 3y_2 \leq 4, \\ & y_1 + y_2 \leq 2, \\ & y_1, y_2 \in \mathbb{R}. \end{aligned}$$

By the complementary slackness, we must have two out of the three constraints binding *at an optimal dual solution y^*.*

(i) *If the second and the third constraints are binding, then we have $y^{(1)} = (2, 0)$, which gives an objective value $w^{(1)} = 5y_1^{(1)} + 8y_2^{(1)} = 10$.*

(ii) *If the first and the third constraints are binding, then we have $y^{(2)} = (1, 1)$, but this is not feasible as it violates the second constraint.*

(iii) *If the first and the second constraints are binding, then we have $y^{(3)} = (-1, 2)$, which gives an objective value $w^{(3)} = 5y_1^{(3)} + 8y_2^{(3)} = 11$.*

Comparing these three possibilities, we see that $y^* = y^{(3)} = (-1, 2)$, which means that $x_3^* = 0$ for any optimal primal solution x^*. This allows us to solve a system of equations for x_1^* and x_2^*, which gives us $x^* = (1, 2, 0)$.

Sometimes it is more efficient to solve the dual problem and use the complementary slackness to recover the primal solution.

Example 4. *Consider the following LO problem*

$$\begin{aligned} \max \quad & 5y_1 + 8y_2 \\ \text{s.t.} \quad & y_1 + 2y_2 \leq 3, \\ & 2y_1 + 3y_2 \leq 4, \\ & y_1 + y_2 \leq 2, \\ & y_2 \leq 1, \\ & y_1, y_2 \in \mathbb{R}. \end{aligned}$$

Note that this is the maximization LO problem in Example 3 with one additional constraint $y_2 \leq 1$. As a result, the solution $y = (-1, 2)$ is no longer feasible for the new maximization problem. Nevertheless, if we write out the dual LO problem

$$\begin{aligned} \min \quad & 3x_1 + 4x_2 + 2x_3 + x_4 \\ \text{s.t.} \quad & x_1 + 2x_2 + x_3 = 5, \\ & 2x_1 + 3x_2 + x_3 + x_4 = 8, \\ & x_1, x_2, x_3, x_4 \geq 0, \end{aligned}$$

we see that the optimal solution to the minimization problem in Example 3 can be extended to a basic feasible solution $x = (1, 2, 0, 0)$ to the new minimization problem. Thus instead of using the big-M or the two-phase simplex method on the maximization LO problem, we can use the known solution to "warm start" our minimization problem. We convert our problem into a maximization problem with the objective function

$$z = -3x_1 - 4x_2 - 2x_3 - x_4,$$

and write the raw tableau as follows.

z	x_1	x_2	x_3	x_4	*rhs*	*basis*
1	3	4	2	1	0	z
0	1	2	1	0	5	x_1
0	2	3	1	1	8	x_2

Through elementary row operations, the standard initial tableau should be as follows.

z	x_1	x_2	x_3	x_4	*rhs*	*basis*
1	0	0	1	-1	-11	z
0	1	0	-1	2	1	x_1
0	0	1	1	-1	2	x_2

By choosing x_4 as the entering variable and x_1 as the leaving variable, we get the new tableau below.

z	x_1	x_2	x_3	x_4	*rhs*	*basis*
1	1/2	0	1/2	0	−21/2	z
0	1/2	0	−1/2	1	1/2	x_4
0	1/2	1	1/2	0	5/2	x_2

This tableau gives an optimal solution $x^* = (0, 5/2, 0, 1/2)$ with an optimal value $z^* = -21/2$. By the complementary slackness, we know that for an optimal dual solution y^*, the second and the fourth constraints must be binding:

$$\begin{aligned} 2y_1^* + 3y_2^* &= 4, \\ y_2^* &= 1. \end{aligned}$$

Then it is straightforward to see that $y^* = (1/2, 1)$, with the objective value $w^* = 5y_1^* + 8y_2^* = 21/2 = -z^*$, which confirms that we have found an optimal solution.

The procedure in Example 4 is a simple illustration of what is known as the *dual simplex method*. The high-level idea behind the dual simplex method is that when a LO problem is modified, we may still be able to reuse some of the found information to make assertions about or accelerate the solution to the modified problem. This idea leads to *sensitivity analysis*, which is discussed below.

3 Sensitivity Analysis

We begin with the assumption that we have found an optimal solution $\bar{x}$ to the primal LO problem (P) with the basis $B \subset \{1, \ldots, n\}$, such that the simplex tableau can be written as

z	x_B	x_N	*rhs*
1	0	$-(c_N - c_B^\top A_B^{-1} A_N)$	$c_B^\top A_B^{-1} b$
0	I	$A_B^{-1} A_N$	$A_B^{-1} b$

(10)

where $N := \{1, \ldots, n\} \setminus B$, and $r = c_N - c_B^\top A_B^{-1} A_N \geq 0$ due to the minimization in (P). We discuss some possible changes to the problem data c, b, and A, respectively, in the following.

3.1 Changing the objective coefficients

We consider a new objective function $z = (c + d)^\top x$ for some $d \in \mathbb{R}^n$, which can be partitioned into d_B and d_N in the same way c is. Note that the solution $\bar{x}$ remains feasible to the modified problem as no constraint is changed. The objective value associated with the solution $\bar{x}$ is changed to

$$(c_B + d_B)^\top A_B^{-1} b = c_B^\top A_B^{-1} b + d_B^\top A_B^{-1} b.$$

The difference here is the product of the change in the basic variable coefficient d_B and the constraint right-hand side $A_B^{-1} b$. To check the optimality of the incumbent solution $\bar{x}$, we need to calculate the modified reduced cost

$$r' := (c_N + d_N) - (c_B + d_B)^\top A_B^{-1} A_N = r + (d_N - d_B^\top A_B^{-1} A_N).$$

The difference here can be calculated by the change of objective coefficients d_B and d_N, and the constraint coefficients in the tableau $A_B^{-1} A_N$. The solution $\bar{x}$ remains optimal if and only if $r' \geq 0$.

3.2 Changing the constraint coefficients

If any constraint coefficient associated with the basic variables A_B is changed, then the inverse A_B^{-1} is also changed (in a nonlinear way), so that we need to recompute it to get the coefficients for x_N and the right-hand sides before we can verify feasibility or optimality of $\bar{x}$. However, if only the coefficients associated with the nonbasic variables A_N are changed, for example, to $A_N + D_N$, then $\bar{x}$ remains feasible with the constraint coefficients $A_B^{-1} A_N$ changed to $A_B^{-1}(A_N + D_N)$. To check the optimality, note that the reduced cost is modified to

$$r' := c_N - c_B^\top A_B^{-1}(A_N + D_N) = r - c_B^\top A_B^{-1} D_N.$$

If we have a dual optimal solution $\bar{y} = (c_B^\top A_B^{-1})^\top$, then the difference in the reduced cost can be calculated more conveniently by $\bar{y}^\top D_N$.

3.3 Changing the constraint right-hand sides

Suppose the constraint right-hand side is changed from b to $b + d$ for some vector $d \in \mathbb{R}^m$. Then the constraint right-hand side is changed to

$$A_B^{-1}(b + d) = A_B^{-1} b + A_B^{-1} d.$$

In general, we would need to compute the inverse A_B^{-1} to obtain the difference $A_B^{-1} d$. If the new right-hand side $A_B^{-1}(b + d)$ is nonnegative, then the solution $\bar{x}$ remains feasible and optimal, because the reduced cost r is unchanged. If the right-hand side $A_B^{-1}(b + d)$ is no longer nonnegative, or if the inverse A_B^{-1} is challenging to compute, while the dual solution $\bar{y}$ is available, then we can use the dual simplex method (as described in

Example 4) to find an optimal dual solution to the problem

$$\begin{aligned} \max \quad & (b+d)^{\mathsf{T}} y \\ \text{s.t.} \quad & A^{\mathsf{T}} y \le c, \\ & y \in \mathbb{R}^m, \end{aligned}$$

and recover a primal optimal solution using complementary slackness. Moreover, if the tableau associated with $\bar{y}$ is known, then the change of the optimal value can be seen from the dual problem with a changed objective coefficient vector.

3.4 Adding a new variable or constraint

If a new variable x_{n+1} is added into our problem (P), then the solution $(x_1, \dots, x_n, x_{n+1}) = (\bar{x}_1, \dots, \bar{x}_n, 0)$ is feasible. We can set B to be the same basis, while $N \leftarrow N \cup \{n+1\}$, so we can start our simplex method from the solution $(\bar{x}, 0)$. Given $a_{n+1} \in \mathbb{R}^n$ and $c_{n+1} \in \mathbb{R}$, we can set

$$D_N = \begin{bmatrix} 0 & \dots & 0 & a_{n+1,1} \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 0 & a_{n+1,m} \end{bmatrix}$$

and calculate the modified reduced cost by

$$r' = (c_N, c_{n+1}) - c_B^{\mathsf{T}} A_B^{-1} D_N.$$

The solution $\bar{x}$ remains optimal if and only if $r' \ge 0$, or equivalently, $c_{n+1} \ge c_B^{\mathsf{T}} A_B^{-1} a_{n+1}$.

If a new constraint is added, then the primal solution $\bar{x}$ may not be feasible any more. However, on the dual side we are basically adding a new variable to the problem (D), so the solution $(\bar{y}, 0)$ remains feasible. Thus we can start our dual simplex method as we did in Example 4.